arXiv:1405.6664v2 [cs.IT] 3 Aug 2014

# On the Computational Intractability of Exact and Approximate Dictionary Learning

Andreas M. Tillmann

***Abstract*—The efficient sparse coding and reconstruction of signal vectors via linear observations has received a tremendous amount of attention over the last decade. In this context, the automated learning of a suitable basis or overcomplete dictionary from training data sets of certain signal classes for use in sparse representations has turned out to be of particular importance regarding practical signal processing applications. Most popular dictionary learning algorithms involve NP-hard sparse recovery problems in each iteration, which may give some indication about the complexity of dictionary learning but does not constitute an actual proof of computational intractability. In this technical note, we show that learning a dictionary with which a given set of training signals can be represented as sparsely as possible is indeed NP-hard. Moreover, we also establish hardness of approximating the solution to within large factors of the optimal sparsity level. Furthermore, we give NP-hardness and non-approximability results for a recent dictionary learning variation called the sensor permutation problem. Along the way, we also obtain a new non-approximability result for the classical sparse recovery problem from compressed sensing.**

***Index Terms*—(SAS-MALN, MLSAS-SPARSE) Machine Learning, Compressed Sensing, Computational Complexity**

## I. Introduction

As a central problem in compressed sensing (CS) [1], [2], [3], the task of finding a sparsest exact or approximate solution to an underdetermined system of linear equations has been a strong focus of research during the past decade. Denoting by $\|\boldsymbol{x}\|_0$ the so-called $\ell_0$-norm, i.e., the number of nonzero entries in $\boldsymbol{x}$, the sparse recovery problem reads

$$\min \|\boldsymbol{x}\|_0 \qquad \text{s.t.} \qquad \|\boldsymbol{D}\boldsymbol{x}-\boldsymbol{y}\|_2 \leq \delta, \qquad (\mathrm{P}_0^\delta)$$

for a given matrix $\boldsymbol{D} \in \mathbb{R}^{m\times n}$ with $m \leq n$ and an estimate $\delta \geq 0$ of the amount of error contained in the measurements $\boldsymbol{y} \in \mathbb{R}^m$. Both the noisefree problem $(\mathrm{P}_0) \coloneqq (\mathrm{P}_0^0)$ and the error-tolerant variant $(\mathrm{P}_0^\delta)$ with $\delta > 0$ are well-known to be NP-hard in the strong sense, cf. [4, problem MP5] and [5], and also difficult to approximate [6], [7].

Groundbreaking results from CS theory include qualificatory conditions (on the dictionary $\boldsymbol{D}$ and the solution sparsity level) which yield efficient solvability of the generally hard problems $(\mathrm{P}_0^\delta)$ by greedy methods—e.g., orthogonal matching pursuit (OMP) [8]—or (convex) relaxations such as Basis Pursuit [9] (replacing the $\ell_0$-norm by the $\ell_1$-norm); see [2], [3], [10] for overviews. Subsequently, numerous optimization algorithms have been tailored to sparse recovery tasks, and various types of dictionaries were shown or designed to exhibit favorable recoverability properties. In particular, the essential assumption of (exact or approximate) sparse representability of certain signal classes using specific dictionaries has been empirically verified in many practical signal processing applications; for instance, natural images are known to admit sparse approximations over discrete cosine or wavelet bases [2].

Nevertheless, a predetermined setup typically cannot fully capture the true structure of real-world signals; thus, using a *fixed* dictionary $\boldsymbol{D}$ naturally restricts the achievable sparsity levels of the representations. Indeed, the simultaneous search for both dictionary and sparse representations of a set of training signals—commonly referred to as *dictionary learning*—was demonstrated to allow for significantly improved sparsity levels using the learned dictionary instead of an analytical, structured or random one. Successful applications of dictionary learning include diverse tasks such as image inpainting [11], [12], denoising [13], [14] and deblurring [15], or audio and speech signal representation [16], [17], to name but a few.

Somewhat informally, the dictionary learning (DL) problem can be stated as: *Given a collection of training data vectors $\boldsymbol{y}^1, \ldots, \boldsymbol{y}^p \in \mathbb{R}^m$ and a positive integer $n$, find a matrix $\boldsymbol{D} \in \mathbb{R}^{m\times n}$ that allows for the sparsest possible representations $\boldsymbol{x}^j$ such that $\boldsymbol{D}\boldsymbol{x}^j = \boldsymbol{y}^j$ (for all $j$).* This task can be formalized in different ways, and there exist many variants seeking dictionaries with further properties such as incoherence [18] or union-of-bases [19]; see also, e.g., [20], [21], [12]. Moreover, several DL algorithms have been developed over the past years; the frequently encountered hard sparse recovery subproblems are typically treated by classical methods from CS. We refer to [22], [23], [24], [25], [26], [12], [27], [28] (and references therein) for a broader overview of well-established DL techniques and some more recent results.

In this paper, we are concerned with the computational complexity of dictionary learning. Due to its combinatorial nature, it is widely believed to be a very challenging problem, but to the best of our knowledge, a formal proof of this intractability claim was missing. We contribute to the theoretical understanding of the problem by providing an NP-hardness proof as well as a strong non-approximability result for DL, see Section II. Furthermore, we prove NP-hardness and non-approximability of an interesting new DL variant—the *sensor permutation problem*, where the sought dictionary is constrained to be related to a given sensing matrix via unknown row permutations; see Section III for the details. As a byproduct, we also obtain a new NP-hardness of approximation result for the sparse recovery problem $(\mathrm{P}_0^\delta)$.

*Remark 1:* Recall that NP-hardness implies that no polynomial-time solution algorithm can exist, under the

A. M. Tillmann is with the Research Group Optimization at TU Darmstadt, Dolivostr. 15, 64293 Darmstadt, Germany (phone: +49-6151-1670868, e-mail: tillmann@mathematik.tu-darmstadt.de).

This work has been accepted by the IEEE for publication. Copyright may be transferred without notice, after which this version may no longer be accessible.

most-widely believed theoretical complexity assumption that P≠NP [4]. Further, *strong* NP-hardness can be understood, in a nutshell, as an indication that a problem's intractability does not depend on ill-conditioning of the input coefficients. This additionally implies that (unless P=NP) there cannot exist a pseudo-polynomial-time exact algorithm and not even a fully polynomial-time approximation scheme (FPTAS), i.e., an algorithm that solves a minimization problem within a factor of $(1+\varepsilon)$ of the optimal value in polynomial time with respect to the input size and $1/\varepsilon$, see [4]. For a thorough and detailed treatment of complexity theory, we refer to [4], [29].

## II. The Complexity of Dictionary Learning

As mentioned in the introduction, different philosophies or goals lead to different formulations of dictionary learning problems, which are usually captured by the general form

$$\min_{\boldsymbol{D},\boldsymbol{X}} f(\boldsymbol{D},\boldsymbol{X};\boldsymbol{Y}) + g(\boldsymbol{D}) + h(\boldsymbol{X}), \tag{1}$$

where the variables are the dictionary $\boldsymbol{D} \in \mathbb{R}^{m\times n}$ (for an a priori chosen $n$) and the matrix $\boldsymbol{X} \in \mathbb{R}^{n\times p}$, whose columns are the representation vectors $\boldsymbol{x}^j$ of the given training signals $\boldsymbol{y}^j$ (w.r.t. the linear model assumption $\boldsymbol{D}\boldsymbol{x}^j \approx \boldsymbol{y}^j$), collected in $\boldsymbol{Y} \in \mathbb{R}^{m\times p}$ as its columns; the functions $f$, $g$ and $h$ express a data fidelity term, and constraints or penalties/regularizers for the dictionary and the representation coefficient vectors, resp.

In the (ideal) noiseless case, the usual approach (see, e.g., [11], [21], [25]) is

$$\min_{\boldsymbol{D},\boldsymbol{X}} \|\boldsymbol{X}\|_0 \quad \text{s.t.} \quad \boldsymbol{D}\boldsymbol{X} = \boldsymbol{Y}, \tag{2}$$

which fits the framework (1) by setting $f(\boldsymbol{D},\boldsymbol{X};\boldsymbol{Y}) := \chi_{\{\boldsymbol{D}\boldsymbol{X}=\boldsymbol{Y}\}}(\boldsymbol{D},\boldsymbol{X})$ (where $\chi$ is the indicator function, i.e., $f(\boldsymbol{D},\boldsymbol{X};\boldsymbol{Y}) = 0$ if $\boldsymbol{D}\boldsymbol{X} = \boldsymbol{Y}$ and $\infty$ otherwise), $g(\boldsymbol{D}) := 0$ and $h(\boldsymbol{X}) := \|\boldsymbol{X}\|_0$ (extending the usual notation to matrices, $\|\boldsymbol{X}\|_0$ counts the nonzero entries in $\boldsymbol{X}$). This problem is a natural extension of $(\mathrm{P}_0)$, and can also be seen as a matrix-factorization problem. To mitigate scaling ambiguities, one often sets $g(\boldsymbol{D}) := \chi_{\{\|\boldsymbol{D}_j\|_2\leq 1\ \forall\, j=1,\dots,n\}}(\boldsymbol{D})$, i.e., the columns of $\boldsymbol{D}$ are required to have bounded norms; cf. [26], [16].

Note that if $n$ is not fixed a priori to a value smaller than $p$, the dictionary learning task becomes trivial: Then, we could just take $\boldsymbol{D} = [\boldsymbol{y}^1,\dots,\boldsymbol{y}^p]$ and exactly represent every $\boldsymbol{y}^i$ using only one column. (Clearly, this also holds for variants which allow representation errors, e.g., $\|\boldsymbol{D}\boldsymbol{x}^j - \boldsymbol{y}^j\|_2 \leq \delta$ for some $\delta > 0$, or minimize such errors under hard sparsity limits $\|\boldsymbol{x}^j\|_0 \leq k$ for some $k \geq 1$.) Thus, requiring $n < p$ is hardly restrictive, in particular since the training data set (and hence, $p$) is usually very large—intuitively, the more samples of a certain signal class are available for learning the dictionary, the better the outcome will be adapted to that signal class—and with respect to storage aspects and efficient (algorithmic) applicability of the learned dictionary, settling for a smaller number of dictionary atoms is well-justified. Similarly, $m \leq n$ is a natural assumption, since sparsity of the coefficient vectors is achieved via appropriate representation bases or redundancy (overcompleteness) of the dictionary; also, at least for large $p$, one can expect $\operatorname{rank}(\boldsymbol{Y}) = m$, in which case $\operatorname{rank}(\boldsymbol{D}) = m \leq n$ becomes necessary to maintain $\boldsymbol{D}\boldsymbol{X} = \boldsymbol{Y}$.

### A. NP-Hardness

As the following results show, finding a dictionary with which the training signals can be represented with optimal sparsity is indeed a computationally intractable problem.

*Theorem 2:* Solving the dictionary learning problem (2) is NP-hard in the strong sense, even when restricting $n = m$.

*Proof:* We reduce from the *matrix sparsification* (MS) problem: Given a full-rankmatrix $\boldsymbol{M} \in \mathbb{Q}^{m\times p}$ $(m < p)$, find a regular matrix $\boldsymbol{B} \in \mathbb{R}^{m\times m}$ such that $\boldsymbol{B}\boldsymbol{M}$ has as few nonzero entries as possible. (The full-rank assumption is not mandatory, but can be made w.l.o.g.: If $\operatorname{rank}(\boldsymbol{M}) = k < m$, $m-k$ rows can be zeroed in polynomial time by elementary row operations, reducing the problem to sparsifying the remaining $k$-row submatrix.) The MS problem was shown to be NP-hard in [30, Theorem 3.2.1] (see also [31], [10]), by a reduction from *simple max cut*, cf. [4, problem ND16]; since this reduction constructs a binary matrix (of dimensions polynomially bounded by the cut problem's input graph size), NP-hardness of MS in fact holds in the strong sense, and we may even assume w.l.o.g. that $\boldsymbol{M} \in \{0,1\}^{m\times p}$.

From an instance of MS, we obtain an equivalent instance of (2) as follows: Set $n := m$ and let $\boldsymbol{Y} := \boldsymbol{M}$. Then, the task (2) is to find $\boldsymbol{D} \in \mathbb{R}^{m\times m}$ and $\boldsymbol{X} \in \mathbb{R}^{m\times p}$ such that $\boldsymbol{D}\boldsymbol{X} = \boldsymbol{Y}$ and $\|\boldsymbol{X}\|_0$ is minimal. (Note that the sought dictionary in fact constitutes a basis for $\mathbb{R}^m$, since $\boldsymbol{M}$ has full row-rank $m$, thus requiring this of the dictionary as well, as discussed above.) Clearly, an optimal solution $(\boldsymbol{D}_*, \boldsymbol{X}_*)$ of this dictionary learning instance gives an optimal solution $\boldsymbol{B}_* = \boldsymbol{D}_*^{-1}$ of MS, with $\boldsymbol{B}_*\boldsymbol{M} = \boldsymbol{X}_*$. It remains to note that the reduction is indeed polynomial, since the matrix inversion can be performed in strongly polynomial time by Gaussian elimination, cf. [32]. Thus, (2) is strongly NP-hard. ∎

*Remark 3:* The above NP-hardness result easily extends to variants of (2) with the additional constraint that, for some constant $c > 0$, $\|\boldsymbol{D}_j\|_2 \leq c$ for all $j$, or $\|\boldsymbol{D}\|_{\mathrm{F}}^2 = \operatorname{tr}(\boldsymbol{D}^\top\boldsymbol{D}) \leq c$ (as treated in [16]): Since the discrete objectives are invariant to scaling in both the dictionary learning and the MS problem, there is always also an optimal $\boldsymbol{D}'_*$ (achieving the same number of nonzeros in the corresponding $\boldsymbol{X}'_*$) that obeys the norm constraints and yields an associated optimal solution $\boldsymbol{B}'_* = (\boldsymbol{D}'_*)^{-1}$ of the MS problem. (Clearly, this argument remains valid for a host of similar norm constraints as well.)

It is not known whether the decision version of the MS problem is contained in NP (and thus not only NP-hard but NP-complete) [30]. Similarly, we do not know if the decision problem associated with (2)—"given $\boldsymbol{Y} \in \mathbb{Q}^{m\times p}$ and positive integers $k$ and $n$, decide whether there exist $\boldsymbol{D} \in \mathbb{R}^{m\times n}$ and $\boldsymbol{X} \in \mathbb{R}^{n\times p}$ such that $\boldsymbol{D}\boldsymbol{X} = \boldsymbol{Y}$ and $\|\boldsymbol{X}\|_0 \leq k$"—is contained in NP, even in the square case $n = m$.

### B. Non-Approximability

Since for NP-hard problems, the existence of efficient (polynomial-time) general exact solution algorithms is deemed impossible, it is natural to search for good approximation methods. Indeed, virtually all well-known dictionary learning algorithms can be interpreted (in a vague sense) as "approximation schemes" since, e.g., the $\ell_0$-norm is convexified to

the $\ell_1$-norm, constraints may be turned to penalty terms in the objective (regularization), etc. However, even disregarding the computational costs of the algorithms, little is known about the quality of the obtained approximations; several recent works along these lines started investigating theoretical recovery properties and error guarantees of dictionary learning algorithms, see, e.g., [33], [34], [28]; in particular, [34] shows the importantance of a good dictionary initialization.

The non-existence of an FPTAS (cf. Remark 1) itself does not generally rule out the existence of an efficient algorithm with some constant approximation guarantee. However, we show below that it is *almost-NP-hard* to approximate the dictionary learning problem (2) to within large factors of the optimal achievable sparsity of representations. Almost-NP-hardness means that no polynomial-time algorithm (here, to achieve the desired approximation ratio) can exist so long as $\mathsf{NP}\not\subseteq\mathsf{DTIME}(N^{\mathrm{poly}(\log N)})$, where $N$ measures the input size (usually, dimension); cf. [35]. This complexity assumption is stronger than $\mathsf{P}\neq\mathsf{NP}$, but also firmly believed (cf., e.g., [36], [37], [38]); it essentially amounts to the claim that not all NP-hard problems admit a quasi-polynomial-time deterministic solution algorithm. Many of the best known non-approximability results are based on this assumption (see, e.g., [39], [36]).

*Theorem 4:* For any $\varepsilon > 0$, the dictionary learning problem (2) cannot be approximated within a factor of $2^{\log^{1-\varepsilon} m}$ in polynomial time, unless $\mathsf{NP}\subseteq\mathsf{DTIME}(m^{\mathrm{poly}(\log m)})$.

*Proof:* In [40], almost-NP-hardness of approximating the optimal value of the matrix sparsification problem[1] to within a factor of $2^{\log^{1/2-o(1)} m}$ (i.e., $2^{\log^{1/2-\varepsilon} m}$ for any $\varepsilon > 0$) was shown, based on results from [35] for the problem of minimizing the number of violated equations in an infeasible linear equation system (see also [6], where this is called MinULR). A closer inspection of [40, Section 3] and [35, Theorems 7 and 8] reveals that this non-approximability result in fact holds up to factors of $2^{\log^{1-\varepsilon} m}$ for any $\varepsilon > 0$. Since our reduction in the proof of Theorem 2 is cost-preserving, this result carries over directly. ∎

This shows that it is extremely unlikely to efficiently learn a dictionary that yields provably good approximations of the sought sparse representations of the training data.

*Remark 5:* The extensions of problem (2) that incorporate norm bounds on $\boldsymbol{D}$ are equally hard to approximate since the respective objectives do not differ from the original matrix sparsification problem's objective, cf. Remark 3. Moreover, the chain of reductions ending in the above result and starting with [35, Theorem 7], maintains a polynomial relationship between the dimensions (here, $m$ and $p$); thus, almost-NP-hardness also holds for approximation to within $2^{\log^{1-\varepsilon} p}$.

*Remark 6:* One may also be interested in learning an *analysis* dictionary $\boldsymbol{\Omega}$, minimizing $\|\boldsymbol{\Omega x}\|_0$ (for given $\boldsymbol{x}$), see, e.g., [41], [42]. Imposing that $\boldsymbol{\Omega}$ has full rank excludes the trivial solution $\boldsymbol{\Omega} = 0$ and, in fact, the square case then is completely equivalent to the MS problem, showing strong NP-hardness and almost-NP-hardness of approximation for analysis dictionary learning; Remarks 3 and 5 apply similarly.

[1]The MS problem in [40] is defined precisely in transposed form compared to the present paper, i.e., there, one seeks to sparsify a full-rank matrix with more rows than columns by right-multiplication with an invertible matrix.

## III. Sparse Coding with Unknown Sensor Locations

Recently, an interesting new problem was introduced in [43] and dubbed the "sensor permutation problem" (for short, SP). Here, it is assumed that the dictionary $\boldsymbol{D}$ is known up to a permutation of its rows, and one wishes to obtain the sparsest representations of the observations $\boldsymbol{Y}$ achievable via permuting these rows—or equivalently, the measurement entries. This approach can model, e.g., faulty wiring in the measurement system setup [43]. Formally, the SP problem can be stated as

$$\min_{\boldsymbol{P},\boldsymbol{X}} \|\boldsymbol{X}\|_0 \quad \text{s.t.} \quad \boldsymbol{AX} = \boldsymbol{PY},\ \boldsymbol{P} \in \mathcal{P}_m, \tag{3}$$

where $\boldsymbol{A} \in \mathbb{R}^{m\times n}$ is a known dictionary, $\boldsymbol{Y} \in \mathbb{R}^{m\times p}$ and $\mathcal{P}_m \coloneqq \{\boldsymbol{P} \in \{0,1\}^{m\times m} : \|\boldsymbol{P}\|_1 = \|\boldsymbol{P}\|_\infty = 1,\ \boldsymbol{P}^\top\boldsymbol{P} = I\}$ denotes the set of all $m \times m$ permutation matrices. (3) can also be seen as a special case of the general dictionary learning framework (1), with $f(\boldsymbol{D},\boldsymbol{X};\boldsymbol{Y}) = \chi_{\{\boldsymbol{DX}=\boldsymbol{Y}\}}(\boldsymbol{D},\boldsymbol{X})$, $g(\boldsymbol{D}) = \chi_{\{\boldsymbol{D}=\boldsymbol{P}^\top\boldsymbol{A}\,:\,\boldsymbol{P}\in\mathcal{P}_m\}}(\boldsymbol{D})$ and $h(\boldsymbol{X}) = \|\boldsymbol{X}\|_0$.

As our following results show, the sensor permutation problem is computationally intractable, even for "nice" input that does not contain numbers of highly varying sizes.

*Theorem 7:* Problem (3) is NP-hard in the strong sense, even if $\boldsymbol{A}$ and $\boldsymbol{Y}$ are binary and $p = 1$. Moreover, for any $\alpha \in (0,1)$ and any $\varepsilon > 0$, there is no polynomial-time algorithm to approximate (3) within a factor of $(1-\alpha)\ln(m)$ unless $\mathsf{P}=\mathsf{NP}$, or to within a factor of $2^{\log^{1-\varepsilon} m}$ unless $\mathsf{NP}\subseteq\mathsf{DTIME}(m^{\mathrm{poly}(\log m)})$. These results remain valid when $\boldsymbol{AX} = \boldsymbol{PY}$ is relaxed to $\|\boldsymbol{AX} - \boldsymbol{PY}\|_2 \leq \delta$ for $0 < \delta \in \mathbb{R}$, and/or $m$ is replaced by $n$.

For the proof, recall the well-known strongly NP-hard Set Cover problem (SC, cf. [4, problem SP5]): "Given a set $\mathcal{S}$ and a collection $\mathcal{C}$ of subsets of $\mathcal{S}$, find a *cover* of minimum cardinality, i.e., a subcollection $\mathcal{C}'$ of as few sets from $\mathcal{C}$ as possible such that $\bigcup_{C\in\mathcal{C}'} C = \mathcal{S}$". A cover $\mathcal{C}'$ is called *exact* if $C \cap D = \emptyset$ for all $C, D \in \mathcal{C}'$ (in other words, if every element of $\mathcal{S}$ is contained in exactly one set from $\mathcal{C}'$).

We will employ the following very recent result:

*Proposition 8 ([44, Theorem 2]):* For every $0 < \alpha < 1$, there exists a polynomial-time reduction from an arbitrary instance of the strongly NP-complete satisfiability problem (SAT, cf. [4, problem LO1]) to an SC instance $(\mathcal{S},\mathcal{C})$ with a parameter $k \in \mathbb{N}$ such that if the input SAT instance is satisfiable, there is an exact cover of size $k$ (and no smaller covers), whereas otherwise, every cover has size at least $(1-\alpha)\ln(|\mathcal{S}|)\,k$.

Recall also that, for any $\varepsilon > 0$, approximating the sparse recovery problem ($\mathrm{P}_0^\delta$) (with any $\delta \geq 0$) to within factors $2^{\log^{1-\varepsilon} n}$ is almost-NP-hard, by [7, Theorem 3]. (In fact, although it clearly goes through for $\delta = 0$ as well, [7] states the proof of this only for $\delta > 0$, because the corresponding result for ($\mathrm{P}_0$) had already been shown in [6] before.) The proof of [7, Theorem 3] is based on a special SC instance construction from [45] (see also [35, Proposition 6]) similar to that from Proposition 8.

*Remark 9:* In the special SC instances underlying the above results, it holds that $|\mathcal{C}|$ and $|\mathcal{S}|$ are polynomially related, so

that all non-approximability results stated in this section also hold with $m$ ($= |\mathcal{S}|$) replaced by $n$ ($= |\mathcal{C}|$).

We are now ready to prove the main result of this section.

*Proof of Theorem 7:* Let $(\mathcal{S}, \mathcal{C}, k, \alpha)$ be a Set Cover instance as in Proposition 8, and let $n = |\mathcal{C}|$, $m = |\mathcal{S}|$. Following the proof of [7, Theorem 3], we first transform the task of finding a minimum-cardinality set cover to the sparse recovery problem ($\mathrm{P}_0$): Define $\boldsymbol{D} \in \{0,1\}^{m \times n}$ by setting $\boldsymbol{D}_{ij} = 1$ if and only if the $i$-th element of $\mathcal{S}$ is contained in the $j$-th set from $\mathcal{C}$, and set $\boldsymbol{y} \coloneqq \mathbb{1}$, i.e., the all-ones vector of length $m$. It is easily seen that the support of every solution $\boldsymbol{x}$ of $\boldsymbol{D}\boldsymbol{x} = \boldsymbol{y}$ induces a set cover (if some element was not covered, at least one row of the equality system would evaluate to $0 = 1$, contradicting $\boldsymbol{D}\boldsymbol{x} = \boldsymbol{y}$). Conversely, every *exact* cover induces a solution of the same $\ell_0$-norm as the cover size (put $\boldsymbol{x}_C = 1$ for the sets $C$ contained in the exact cover, and zero in the remaining components). Thus, if there is an *exact* cover of size $k$, there is a $k$-sparse solution of $\boldsymbol{D}\boldsymbol{x} = \boldsymbol{y}$. Conversely, if all set covers have size at least $(1-\alpha)\ln(m)\,k$, then necessarily all $\boldsymbol{x}$ with $\boldsymbol{D}\boldsymbol{x} = \boldsymbol{y}$ have $\|\boldsymbol{x}\|_0 \geq (1-\alpha)\ln(m)\,k$ (because otherwise, the support of $\boldsymbol{x}$ would yield a set cover of size smaller than $(1-\alpha)\ln(m)\,k$).

This instance of ($\mathrm{P}_0$) is now easily transformed into one of the sensor permutation problem (3): We set $\boldsymbol{A} \coloneqq \boldsymbol{D}$, $\boldsymbol{Y} \coloneqq \boldsymbol{y}$ (thus, $p = 1$). Now, since $\boldsymbol{Y} = \mathbb{1}$, $\boldsymbol{P}\boldsymbol{Y} = \boldsymbol{Y}$ for all $\boldsymbol{P} \in \mathcal{P}_m$ and the choice of $\boldsymbol{P}$ has no influence on the solution. Thus, indeed, the SP problem (3) for these $\boldsymbol{A}$ and $\boldsymbol{Y}$ has precisely the same solution value as the above-constructed instance of ($\mathrm{P}_0$). Since solving the original Set Cover instance is (strongly) NP-hard (by Proposition 8), and all constructed numbers and their encoding lengths remain polynomially bounded by the input parameter $m$ (and $n$), this immediately shows the claimed strong NP-hardness result. In fact, could we approximate, in polynomial time, the optimal solution value of (3) to within a factor of $(1-\alpha)\ln(m)$, then we could also decide the SAT instance underlying the SC problem from Proposition 8 in polynomial time, which is impossible unless P=NP. Therefore, for any $0 < \alpha < 1$, even approximating (3) to within factors $(1-\alpha)\ln(m)$ is NP-hard.

For the second non-approximability result of Theorem 7, it suffices to note that the construction above is cost-preserving and that the ($\mathrm{P}_0$) instance in the proof of [7, Theorem 3] also has $\boldsymbol{y} = \mathbb{1}$. Hence, we can directly transfer the non-approximability properties, and conclude that there is no polynomial-time algorithm approximating (3) to within factors $2^{\log^{1-\varepsilon} n}$ (for any $\varepsilon > 0$), unless $\mathsf{NP} \subseteq \mathsf{DTIME}(n^{\mathrm{poly}(\log n)})$.

Finally, the above results extend to the noise-aware SP problem variant by treating the relaxed constraints $\|\boldsymbol{A}\boldsymbol{X} - \boldsymbol{P}\boldsymbol{Y}\|_2 \leq \delta$ for $\delta > 0$ completely analogously to the proof of [7, Theorem 3] (we omit the details) and, by Remark 9, remain valid w.r.t. either $m$ or $n$. ∎

*Remark 10:* The decision version of (3) is easily seen to be in NP (for rational input), and hence NP-complete.

Note that the first part of the above proof yields a new non-approximability result for sparse recovery:

*Corollary 11:* For any $\alpha \in (0,1)$, it is NP-hard to approximate ($\mathrm{P}_0^\delta$) to within a factor of $(1-\alpha)\ln(n)$.

This complements the previously known results from [6, Theorem 7] and [7, Theorem 3]: For $n$ large enough (and some fixed pair $\alpha$, $\varepsilon$), $2^{\log^{1-\varepsilon} n} > (1-\alpha)\ln(n)$, but the assumption P≠NP is weaker than $\mathsf{NP} \not\subseteq \mathsf{DTIME}(n^{\mathrm{poly}(\log n)})$.

## IV. Concluding Remarks

In this note, we gave formal proofs for NP-hardness and non-approximability of several dictionary learning problems. While perhaps not very surprising, these results provide a complexity-theoretical justification for the common approaches to tackle dictionary learning tasks by inexact methods and heuristics without performance guarantees.

While preparing this manuscript, we became aware of a related result presented at ICASSP 2014, see [46]. In that work, the authors claim NP-hardness of approximating

$$\min_{\boldsymbol{D},\boldsymbol{X}} \|\boldsymbol{D}\boldsymbol{X} - \boldsymbol{Y}\|_{\mathrm{F}}^2 \quad \text{s.t.} \quad \|\boldsymbol{X}_j\|_0 \leq k \;\; \forall\, j = 1, \ldots, p, \quad (4)$$

to within a given *additive* error w.r.t. the objective (i.e., not within a *factor* of the optimal value), for the case in which $\boldsymbol{Y}$ contains only two columns and $k$ is fixed to 1. Unfortunately, [46] does not contain a proof, and at the time of writing, we could not locate it elsewhere. Note also that, clearly, (4) is also a special case of the general formulation (1)—using $f(\boldsymbol{D}, \boldsymbol{X}; \boldsymbol{Y}) = \|\boldsymbol{D}\boldsymbol{X} - \boldsymbol{Y}\|_{\mathrm{F}}^2$, $h(\boldsymbol{X}) = \chi_{\{\|\boldsymbol{X}_j\|_0 \leq k \;\forall j\}}(\boldsymbol{X})$ and $g(\boldsymbol{D}) = 0$—but that the results from the present paper and from [46] nevertheless pertain to different problems, both of which are often referred to as "dictionary learning".

Future research closely related to the present work could include investigating the potential use of matrix sparsification based heuristics for dictionary learning purposes (e.g., when learning a union-of-bases dictionary as in [19]).

Note also that the reduction from [40] does not admit transferring the NP-hardness of approximating MinULR to within any constant factor (see [35, Theorem 5]) to the MS problem. (Similarly, the reduction to MS in [30] apparently does not preserve approximation ratios.) Such non-approximability results under the slightly weaker P≠NP assumption hence remain open for problem (2) (and its norm-constrained variants). Also, the complexities of dictionary learning with $\ell_1$-objective and/or noise-awareness (e.g., constraints $\|\boldsymbol{D}\boldsymbol{X} - \boldsymbol{Y}\|_{\mathrm{F}} \leq \delta$ for $\delta > 0$) remain important open problems.

On the other hand, one may wish to focus on "good news", e.g., by designing efficient approximation algorithms that give performance guarantees not too much worse than our intractability thresholds, or by identifying special cases which are notably easier to solve. Also, it would be interesting to develop further "hybrid algorithms" that combine relaxation methods and tools from combinatorial optimization, such as the branch & bound procedure from [43].

## Acknowledgments

The author would like to thank Yonina Eldar and Julien Mairal for enticing him to look into dictionary learning, Rémi Gribonval for bringing the sensor permutation problem to his attention, as well as Imke Joormann, Marc Pfetsch and two anonymous referees for their valuable comments on an earlier version of the manuscript.

## REFERENCES


[1] D. L. Donoho, "Compressed Sensing," *IEEE Trans. Inform. Theory*, vol. 52, no. 4, pp. 1289–1306, 2006.

[2] S. Foucart and H. Rauhut, *A Mathematical Introduction to Compressive Sensing*, ser. Applied and Numerical Harmonic Analysis. Birkhäuser, 2013.

[3] G. Kutyniok and Y. C. Eldar, Eds., *Compressed Sensing: Theory and Applications*. Cambridge University Press, 2012.

[4] M. R. Garey and D. S. Johnson, *Computers and Intractability. A Guide to the Theory of NP-completeness*. W. H. Freeman and Company, 1979.

[5] B. K. Natarajan, "Sparse Approximate Solutions to Linear Systems," *SIAM J. Comput.*, vol. 24, no. 2, pp. 227–234, 1995.

[6] E. Amaldi and V. Kann, "On the Approximability of Minimizing Nonzero Variables or Unsatisfied Relations in Linear Systems," *Theor. Comput. Sci.*, vol. 209, no. 1–2, pp. 237–260, 1998.

[7] E. Amaldi, "On the complexity of designing compact perceptrons and some consequences," in *El. Proc. 5th Internat. Symp. on Artificial Intelligence and Math.*, 1999.

[8] Y. C. Pati, R. Rezaiifar, and P. S. Krishnaprasad, "Orthogonal Matching Pursuit: Recursive Function Approximation with Applications to Wavelet Decomposition," in *Proc. 27th Ann. Asilomar Conference on Signals, Systems and Computers*. IEEE Computer Society Press, 1993, vol. 1, pp. 40–44.

[9] S. S. Chen, D. L. Donoho, and M. A. Saunders, "Atomic Decomposition by Basis Pursuit," *SIAM J. Sci. Comput.*, vol. 20, no. 1, pp. 33–61, 1998.

[10] A. M. Tillmann, "Computational Aspects of Compressed Sensing," Doctoral dissertation, TU Darmstadt, Germany, 2013.

[11] M. Aharon, M. Elad, and A. M. Bruckstein, "On the uniqueness of overcomplete dictionaries, and a practical way to retrieve them," *Linear Algebra Appl.*, vol. 416, no. 1, pp. 48–67, 2006.

[12] J. Mairal, F. Bach, J. Ponce, and G. Shapiro, "Online Learning for Matrix Factorization and Sparse Coding," *J. Mach. Learn. Res.*, vol. 11, pp. 19–60, 2010.

[13] M. Elad and M. Aharon, "Image Denoising Via Sparse and Redundant Representations Over Learned Dictionaries," *IEEE Trans. Image Process.*, vol. 15, no. 12, pp. 3736–3745, 2006.

[14] S. Beckouche, J. L. Starck, and J. Fadili, "Astronomical image denoising using dictionary learning," *Astron. Astrophys.*, vol. 556, no. A132, 2013.

[15] F. Couzinie-Devy, J. Mairal, F. Bach, and J. Ponce, "Dictionary Learning for Deblurring and Digital Zoom," arXiv:1110.0957 [cs.LG], 2011.

[16] M. Yaghoobi, T. Blumensath, and M. Davies, "Regularized Dictionary Learning for Sparse Approximation," in *Proc. EUSIPCO'08*, 2008.

[17] M. G. Jafari and M. D. Plumbley, "Fast Dictionary Learning for Sparse Representation of Speech Signals," *IEEE J. Sel. Top. Signa.*, vol. 5, no. 5, pp. 1025–1031, 2011.

[18] D. Barchiesi and M. D. Plumbley, "Learning Incoherent Dictionaries for Sparse Approximation Using Iterative Projections and Rotations," *IEEE Trans. Signal Process.*, vol. 61, no. 8, pp. 2055–2065, 2013.

[19] S. Lesage, R. Gribonval, F. Bimbot, and L. Benaroya, "Learning unions of orthonormal bases with thresholded singular value decompositon," in *Proc. IEEE ICASSP'05*, 2005, vol. 5, pp. v/293–v/296.

[20] R. Rubinstein, M. Zibulevsky, and M. Elad, "Double Sparsity: Learning Sparse Dictionaries for Sparse Signal Approximation," *IEEE Trans. Signal Process.*, vol. 58, no. 3, pp. 1553–1564, 2010.

[21] M. D. Plumbley, "Dictionary Learning for L1-Exact Sparse Coding," in *Independent Component Analysis and Signal Separation (Proc. ICA'07)*, ser. Lect. Notes Comput. Sc. Springer, 2007, vol. 4666, pp. 406–413.

[22] K. Engan, S. O. Aase, and J. H. Husøy, "Method of Optimal Directions for Frame Design," in *Proc. IEEE ICASSP'99*, 1999, vol. 5, pp. 2443–2446.

[23] M. Aharon, "Overcomplete Dictionaries for Sparse Representation of Signals," Ph.D. dissertation, Technion – Israel Institute of Technology, Haifa, Israel, 2006.

[24] M. Aharon, M. Elad, and A. M. Bruckstein, "K-SVD: An Algorithm for Designing of Overcomplete Dictionaries for Sparse Representations," *IEEE Trans. Signal Process.*, vol. 54, no. 11, pp. 4311–4322, 2006.

[25] R. Gribonval and K. Schnass, "Dictionary Identification – Sparse Matrix-Factorization via $\ell_1$-Minimization," *IEEE Trans. Inform. Theory*, vol. 56, no. 7, pp. 3523–3539, 2010.

[26] K. Kreutz-Delgado, J. F. Murray, B. D. Rao, K. Engan, T.-W. Lee, and T. J. Sejnowski, "Dictionary Learning Algorithms for Sparse Representation," *Neural Comput.*, vol. 15, no. 2, pp. 349–396, 2003.

[27] A. Rakotomamonjy, "Direct Optimization of the Dictionary Learning Problem," *IEEE Trans. Signal Process.*, vol. 61, no. 22, pp. 5495–5506, 2013.

[28] S. Arora, A. Bhaskara, R. Ge, and T.Ma, "More Algorithms for Provable Dictionary Learning," arXiv:1401.0579 [cs.DS], 2014.

[29] B. Korte and J. Vygen, *Combinatorial Optimization. Theory and Algorithms*, 5th ed., ser. Algorithms and Combinatorics. Springer, 2011, vol. 21.

[30] S. T. McCormick, "A Combinatorial Approach to some Sparse Matrix Problems," Ph.D. dissertation, Stanford University, CA, USA, 1983.

[31] T. F. Coleman and A. Pothen, "The Sparse Null Space Basis Problem," Cornell University, Ithaca, NY, USA, Tech. Rep. TR 84-598, 1984.

[32] M. Grötschel, L. Lovász, and A. Schrijver, *Geometric Algorithms and Combinatorial Optimization*, 2nd ed., ser. Algorithms and Combinatorics. Springer, 1993, vol. 2.

[33] D. A. Spielman, H. Wang, and J. Wright, "Exact Recovery of Sparsely-Used Dictionaries," *J. Mach. Learn. Res.*, vol. 23, pp. 37.1–37.18, 2012.

[34] A. Agarwal, A. Anandkumar, P. Jain, P. Netrapalli, and R. Tandon, "Learning Sparsely Used Overcomplete Dictionaries via Alternating Minimization," arXiv:1310.7991 [cs.LG], 2013.

[35] S. Arora, L. Babai, J. Stern, and Z. Sweedyk, "The Hardness of Approximate Optima in Lattices, Codes, and Systems of Linear Equations," *J. Comput. Syst. Sci.*, vol. 54, no. 2, pp. 317–331, 1997.

[36] L. Trevisan, "Inapproximability of Combinatorial Optimization Problems," in *Paradigms of Combinatorial Optimizations: Problems and New Approaches*, V. T. Paschos, Ed. John Wiley & Sons, 2013, vol. 2, pp. 381–434.

[37] H. Buhrman and S. Homer, "Superpolynomial Circuits, Almost Sparse Oracles and the Exponential Hierarchy," in *Proc. FSTTCS 12*, ser. Lect. Notes Comput. Sci. Springer, 1992, vol. 652, pp. 116–127.

[38] R. Impagliazzo and R. Paturi, "On the Complexity of $k$-SAT," *J. Comput. Syst. Sci.*, vol. 62, pp. 367–375, 2001.

[39] S. Arora and C. Lund, "Hardness of Approximations," in *Approximation Algorithms for NP-hard Problems*, D. Hochbaum, Ed. PWS Publishing, 1996, pp. 399–446.

[40] L.-A. Gottlieb and T. Neylon, "Matrix Sparsification and the Sparse Null Space Problem," in *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (Proc. APPROX'10 and RANDOM'10)*, ser. Lect. Notes Comput. Sci. Springer, 2010, vol. 6302, pp. 205–218.

[41] R. Rubinstein, T. Peleg, and M. Elad, "Analysis K-SVD: A Dictionary-Learning Algorithm for the Analysis Sparse Model," *IEEE Trans. Signal Process.*, vol. 61, no. 3, pp. 661–677, 2013.

[42] M. Yaghoobi, S. Nam, R. Gribonval, and M. E. Davies, "Analysis Operator Learning for Overcomplete Cosparse Representations," in *Proc. EUSIPCO'11*, 2011.

[43] V. Emiya, A. Bonnefoy, L. Daudet, and R. Gribonval, "Compressed Sensing with Unknown Sensor Permutation," in *Proc. IEEE ICASSP'14*, 2014, pp. 1040–1044.

[44] D. Moshkovitz, "The Projection Games Conjecture and the NP-Hardness of $\ln n$-Approximating Set-Cover," Preprint, 2014. [Online]. Available: http://people.csail.mit.edu/dmoshkov/papers/set-cover/set-cover-full.pdf

[45] M. Bellare, S. Goldwasser, C. Lung, and A. Russell, "Efficient Probabilistically Checkable Proofs with Applications to Approximation Problems," in *Proc. 25th ACM Symp. Theory Comput.*, 1993, pp. 294–304.

[46] M. Razaviyayn, H.-W. Tseng, and Z.-Q. Luo, "Dictionary Learning for Sparse Representation: Complexity and Algorithms," in *Proc. IEEE ICASSP'14*, 2014, pp. 5247–5251.